11. The apparatus of claim 10, wherein the means for calculating action-values for a plurality of actions based on a current state using a trained deep neural network comprises:

means for extracting features from the first medical image using first convolutional layers of the trained deep neural network;

means for extracting features from the second medical image using second convolutional layers of the trained deep neural network; and

means for comparing the extracted features from the first medical image and the extracted features from the second medical image to calculate the action-values for the plurality of actions.

12. The apparatus of claim 11, wherein the first convolutional layers and the second convolutional layers are implemented with separate weights.

13. The apparatus of claim 10, wherein the means for computing a deformation field based on the selected action and a deformation model comprises:

means for encoding deformation for an entirety of the first medical image using a dense model.

14. The apparatus of claim 13, wherein the means for applying the deformation field to the first medical image to warp the first medical image comprises:

means for decoding the deformation for the entirety of the first medical image; and

means for applying the decoded deformation to the first medical image.

15. A non-transitory computer readable medium storing computer program instructions for deformable registration of medical images using an intelligent artificial agent, the computer program instructions defining operations comprising:

receiving a first medical image and a second medical image;

calculating action-values for a plurality of actions based on a current state using a trained deep neural network, the current state determined from the first medical image and the second medical image;

selecting an action from the plurality of actions based on the calculated action-values;

computing a deformation field based on the selected action and a deformation model;

applying the deformation field to the first medical image to warp the first medical image; and

repeating the calculating, the selecting, the computing, and the applying for a plurality of iterations using the warped first medical image as the first medical image to register the first medical image and the second medical image.

16. The non-transitory computer readable medium of claim 15, wherein calculating action-values for a plurality of actions based on a current state using a trained deep neural network comprises:

extracting features from the first medical image using first convolutional layers of the trained deep neural network;

extracting features from the second medical image using second convolutional layers of the trained deep neural network; and

comparing the extracted features from the first medical image and the extracted features from the second medical image to calculate the action-values for the plurality of actions.

17. The non-transitory computer readable medium of claim 15, wherein computing a deformation field based on the selected action and a deformation model comprises:

encoding deformation at particular points in the first medical image that have corresponding points in the second medical image using a correspondence based model.

18. The non-transitory computer readable medium of claim 17, wherein applying the deformation field to the first medical image to warp the first medical image comprises:

decoding the deformation at the particular points in the first medical image;

transforming a mesh of a segmented structure in the first medical image based on the decoded deformation;

computing a dense deformation field based on the transformation of the mesh; and

applying the computed dense deformation field to the first medical image.

19. The non-transitory computer readable medium of claim 15, wherein computing a deformation field based on the selected action and a deformation model comprises:

modeling deformation of pixels at a location in the first medical image as having a Gaussian distribution; and

computing a deformation at a particular point in the first medical image based on a distance between the particular point and the location, a magnitude of the Gaussian distribution, and a standard deviation of the Gaussian distribution.

20. The non-transitory computer readable medium of claim 15, wherein the trained deep neural network is trained to predict the action-values for the plurality of actions based on a plurality of training image pairs with known ground truth transformation parameters using reinforcement learning in which, for each training image pair, a reward for each action of the plurality of actions at a given state is used to train the deep neural network to learn an optimal registration policy.

* * * * *